

What is claimed is:

1. A memory storage device having an operating system which uses at least one
inode for accessing file segments, the inode comprising:
a plurality of rows; and
5 a portion of the rows storing extents pointing to data blocks, each extent having a
field to indicate whether the extent is an indirect extent, a hole extent or a direct extent.
2. The memory storage device of claim 1, wherein
each inode is adapted to allow any portion of the extents stored therein to be indirect
10 extents.
3. The memory device of claim 1, wherein a portion of the extents are hole extents.
4. The memory device of claim 1, wherein each extent further comprises a length
15 field, the length field of each indirect extent indicating the number of data blocks
pointed to indirectly by the indirect extent.
5. An automated method of storing data files in a memory storage system,
comprising:
20 assigning an inode to a data file to be stored; and
writing a plurality of extents in the inode, each extent pointing to a string of one
or more data blocks for storing a segment of the data file and having a field for indicating
that the extent is one of an indirect extent, a hole extent, and a direct extent.
- 25 6. The method of claim 5, further comprising:

replacing each of a plurality of the direct extents by at least one indirect extent pointing to a data block; and

writing to each data block pointed to by one of the indirect extents the direct extent that is replaced by the one of the indirect extents.

5

7. A method of storing data files, comprising:

writing a plurality of extents to an inode;

writing first and second data segments of the file to first and second physical data blocks, first and second ones of the extents pointing to the first and second

10 physical data blocks;

shifting at least one of the extents in the inode to leave a space between the first and second ones of the extents in response to the inode having at least one empty row; and

15 inserting a third direct extent in the space in response to the inode having an empty row between the first and second extents, the direct extent pointing to a third data block to store a new data segment of the file.

8. The method of claim 7, further comprising:

20 writing the third direct extent and the second extent to an indirect block in response to the inode not having an empty row between the first and second extents; and

inserting an indirect extent pointing to the indirect block to the row of the inode occupied by the second extent.

9. The method of claim 8, further comprising:

25 updating a length field of the indirect extent to include a number of blocks indirectly pointed to by the indirect extent.

10. In a memory storage device employing an operating system comprising instructions for a method for storing data files, the method comprising:

writing a plurality of extents to an inode assigned 5 to a file, including first and second direct extents;

5 writing data to first and second data blocks, said first and second direct extents pointing to the first and second physical data blocks;

inserting an indirect extent in the inode between the first and second direct extents, the indirect extent pointing to a third physical data block; and

10 writing at least one extent to the third physical data block, the one extent pointing to a physical data block storing a segment of the file.

11. The method of claim 10, further comprising:

writing a length field to each extent, the length field fixing the number of consecutive data blocks pointed to by the extent.

15

12. The method of claim 10, the operating system being a UNIX based system.

13. A distributed storage system, comprising:

a global cache memory;

20 a plurality of processors coupled to the global cache memory, each processor having a local memory for storing an operating system; and

a plurality of data storage devices coupled to the global cache memory, the devices and processors capable of communicating by posting messages to each other in the cache memory, each of the devices having a processor and local memory storing an
25 operating system, each operating system including an extent based file system for abstracting file names to physical data blocks in the storage devices, wherein each extent

includes a field to indicate whether the extent points to a block of extents or a block of data.

5 14. The system of claim 13, wherein each operating system is adapted to map files to data blocks by assigning an inode to a file, each inode capable of storing a plurality of extents.

15. The system of claim 14, wherein each extent includes a field to indicate whether the extent points to a block of extents or a block of data.

10

16. The system of claim 13, each operating system being a UNIX based system.